

REMARKS

Status of the Claims

Claims 1 - 36 are now pending in the present application, Claims 14 and 33 having been amended to correct typographical errors and Claims 1, 19, 24, and 36 having been amended to more clearly distinguish the subject matter from the cited art.

Telephone Interview Request

In the event that the Examiner is not persuaded by the following arguments that these claims are in condition for allowance, applicants respectfully request that the Examiner telephone applicants' attorney to schedule a mutually convenient time for a Telephone Interview in order to discuss how applicants might advance prosecution of the above-identified application to issue.

Claims Rejected Under 35 U.S.C. § 103(a)

The Examiner has rejected Claims 1, 2, 4-8, 10-27, 29-34, and 36 under 35 U.S.C. § 103(a) as being unpatentable by U.S. Patent No. 6,128,279 (O’Neil et al., hereinafter referred to as “O’Neil”) in view of U.S. Patent No. 6,438,652 (Jordan et al., hereinafter referred to as “Jordan”).

In the interest of reducing the complexity of the issues for the Examiner to consider in this response, the following discussion focuses on amended independent Claims 1, 19, 24, and 36. The patentability of each dependent claim is not necessarily separately addressed in detail. However, applicants' decision not to discuss the differences between the cited art and each dependent claim should not be considered as an admission that applicants concur with the Examiner's conclusion that these dependent claims are not patentable over the cited references. Similarly, applicants' decision not to discuss differences between the prior art and every claim element, or every comment made by the Examiner, should not be considered as an admission that applicants concur with the Examiner's interpretation and assertions regarding those claims. Indeed, applicants believe that all of the dependent claims patentably distinguish over the references cited. However, a specific traverse of the rejection of each dependent claim is not required, since dependent claims are patentable for at least the same reasons as the independent claims from which the dependent claims ultimately depend.

Rejection of Independent Claim 1

Significant differences exist between applicants' Claim 1 and the cited combination of O'Neil and Jordan because the combined references do not teach or suggest the functionality of applicants' intake

1 message as recited in step (d) namely, that the intake message identifies the second resource as an
2 *intake*, and specifies that “*the new intake is a single entry-point service resource in the cluster.*”

3 Applicants’ step (d) as amended recites (with emphasis added):

4 designating the second resource as the new intake and providing an intake
5 message from the first resource to the plurality of resources in the cluster *identifying*
6 *the second resource as the new intake, wherein the new intake is a single entry-point*
service resource in the cluster;

7 The Examiner has asserted that the cited art teaches applicants’ step (d). Although O’Neil
8 does not specifically teach that the first intake provides a message to the cluster identifying the
9 second resource as the intake, the Examiner asserts that Jordan teaches the communication of a server
10 that indicates that it is no longer available to process requests, and requests the routing to another
11 server, which the Examiner indicates constitutes the announcement of intake designation to a second
12 server. In support of his assertion, the Examiner cites column 4, lines 15-27, which is reproduced
13 below:

14 In another example, an overloaded cooperating cache server can identify a less
15 loaded cooperating cache server; and communicate a shift request and a copy of the cached
16 object to the less loaded cooperating cache server (which then caches the object), so that
17 subsequent requests for the object will not be forwarded. Alternatively, an overloaded
18 cooperating cache server can communicate the shift request to the less loaded cooperating
19 cache server, which then obtains a copy of the object from an originating object server, in
20 response to the shift request. In yet another alternative, the owning cache server can
multicast the shift request message to one or more of the other cooperating cache servers so
that subsequent forward requests will be shifted. (Emphasis added, Jordan, column 4,
lines 14-27.)

21
22 The Examiner concludes that it would have been obvious to one of ordinary skill in the art at the
23 time of the invention to include the communication of a message, to indicate that a certain server is no
24 longer the intake server, as taught by Jordan in the system of O’Neil. The Examiner asserts that the
25 motivation for doing so is to enable a server in use to communicate with a system that it is unavailable for
26 further requests, to thus provide more efficiency in the system, by effectively disallowing overburdening the
27 server. The Examiner further notes that both the present application and the inventions discussed in the
28 cited art are directed to the same field of endeavor, namely the load balancing of servers.

1 Jordan Does Not Teach Or Suggest An Equivalent of Applicants' "Intake"

2 Under the section entitled "Response to Arguments" the Examiner cites the underlined
3 portion of the above citation as teaching that the shift request identifies a less-loaded server, as
4 further evidenced by the use of the word "the" in the cited portion, in describing a less-loaded server.
5 Therefore, the Examiner concludes that this constitutes identification of the server that will serve as
6 the new intake. In addition, the Examiner asserts that a message is sent to the other servers,
7 indicating that the less-loaded server has been chosen as the one servicing new requests. It is further
8 asserted that Jordan thus provides the teaching of an intake message, designating a new and specific
9 server as the new intake, sent out by the original intake. Further, the Examiner asserts that a message
10 is sent to the other servers, indicating that the less-loaded server has been chosen as the one servicing
11 new requests, which the Examiner asserts constitutes the teaching of an intake message, designating a
12 new and specific server as the new intake, sent out by the original intake. The Examiner next
13 indicates that further support for this teaching may be found in column 9, lines 9-15 and 33-36 of
14 Jordan, wherein the multicasting of a message indicating that further requests will be routed to a new
15 specific owner is taught. The Examiner also states that a message indicating a new shared ownership
16 is taught, which in the event of a shift request message by a first shared owner, new requests will be
17 routed to the second shared owner, and not the first one, which further constitutes another example of
18 Jordan's disclosure of an intake message.

19 Applicants have amended step (d) in Claim 1 to clarify that the second resource is a new
20 *intake*, wherein an intake is defined as *a single entry-point service resource in a cluster* (see
21 applicants' specification, page 7, lines 13-14). A separate intake is designated for each different type
22 of service that is being provided by the cluster such that clients are grouped together and processed in
23 a group at a service resource for as long as the service is provided by the resource (see applicants'
24 specification, page 7, lines 14-18). Applicants' specification provides:

25 One game service instance of each type of game is designated as an "intake"
26 for the cluster. The intake for a specific type of game serves as a central point for new
client requests and groups, or matches, clients into a game. For example, Checkers
game service 112aa is designated as the Checkers intake 124 for the Checkers game
type. Similarly, Hearts game service 110bb is designated as the Hearts intake 126 for
the Hearts game type. Any game service instance can be designated the intake for that
27 game type regardless of which node within the cluster is executing the designated
28 game service instance. However, only one game service instance can be designated as
29 the intake for each type of game. A game service instance that is designated as the
30

1 intake for a game type is responsible for accepting new client requests to participate in
2 that type of game. (Emphasis added, applicants' specification, page 13, line 36-
3 page 14, line 9).

4 For example, as disclosed in the underlined portion above, an intake is a game service
5 instance of each type of game that is designated such as a Checkers intake 124, while another is
6 designated as a Hearts intake 126 and these intakes will process the group for as long as the service is
7 provided. In contrast, Jordan does not disclose any equivalent of an intake meeting the recited
8 definition of *a single entry-point service resource in a cluster*, but instead, discloses an owning cache
9 server (see Jordan, column 9, line 31) as illustrated in Figure 2a under ownership 1012. Jordan's
10 Figure 4 shows an example of a logic flow chart for a cache server in response to a request for an
11 object either directly 155 from a browser 160 or forwarded 125 from the load monitor 120. Step 305
12 describes how the cache server (that received the request) will fetch the object from the originating
13 web server and return the object (Jordan, column 7, lines 33-36). The process then ends, in a
14 step 306. Thus, Jordan's servers are simply cache servers that fetch and return objects and are
15 therefore not equivalent to an intake that is a single entry-point service resource in a cluster, as
16 recited by applicants' claim.

17 With respect to the Examiner's assertion that Jordan teaches a message that identifies a less
18 loaded server, applicants respectfully disagree and note that Jordan's *shift request* does not identify a
19 less-loaded server that will serve as the new intake, as is evident from the following:

20 For example, the present invention can be configured to perform load
21 balancing for a collection of cooperating proxy cache servers where each cache
22 server 150 multicasts to a list of cooperating cache servers to locate a copy of a locally
23 missed object. In this case, no specific ownership information need be maintained
24 anywhere in the system. However, there is also no guarantee of finding an object from
25 the cooperating cache servers, either. Assume that a logical load monitor 120 is used
26 to maintain the load conditions 1021 of all proxy cache servers and share this
27 information with each cache server 150. The load balancing can be achieved by
28 excluding overloaded servers from the list of cooperating servers to which a cache
29 server multicasts its request (also called a shift request). As a result, only less loaded
30 cache servers will receive forwarded requests 125. (Emphasis added, Jordan,
 column 7, line 66-column 8, line 13).

31 As stated in the underlined portions of the above quote, Jordan's shift request is a request
32 wherein a cache server multicasts to a **list** of other cache servers in order to locate a copy of a locally

1 missed object. Overloaded servers are excluded from the list to which the cache server multicasts the
2 request, which Jordan indicates is also called a *shift request*. The shift request is thus sent to those
3 servers on a list that are not overloaded, but a shift request is not an intake message. The shift request
4 is multicast to *a list of servers that are not overloaded*, but that step is not equivalent to identifying *a*
5 *specific server* that is not overloaded and has been designated as the new intake. Thus, Jordan does
6 not teach an equivalent of applicants' intake message as recited in subparagraph (d) of Claim 1.

7 With respect to the Examiner's assertion regarding a message being sent to the other servers,
8 Jordan does NOT teach that a message is sent to the other servers, indicating that a specific less-
9 loaded server has been chosen to service new requests, identifying that server as being the new
10 intake. Jordan's Figure 3 shows an example of a logic flow chart for the load monitor in response to
11 a request from a cache server because of a cache miss and its accompanying disclosure is reproduced
12 below:

13 FIG. 3 shows an example of a logic flow for steps taken by the load monitor
14 120 in response to a request 125 from a cache server 150 because of a cache miss. As
15 depicted, in step 201, it checks to see if the requested object/partition can be found in
16 the caching table. If not, in step 202, a new entry is created for the object/partition and
17 a cache server is assigned as its owner. After the entry is located in the caching table,
18 in step 203, the forwarding frequency 1011 is updated, e.g., incremented by 1. The
19 load monitor then examines the load table 102 to see if the owner is currently
20 overloaded (and that the forwarding frequency 1011 is a significant contributor
21 thereto), in step 204. If yes, in step 205, the load monitor finds an underloaded (or less
22 loaded) cache server and assign it as the new 10122 (or shared) owner 10122 of the
23 requested object. The ownership information 1012 for the object in the caching table
24 101 is updated accordingly. Those skilled in the art will appreciate that the logic flow
25 could comprise a shared 10123 or hierarchical ownership 1012 in the caching table
26 101 or other data structure employed. The request (possibly with a copy of the
27 requested object) can then be forwarded 125 to a new sole 10122 (or shared 10123)
28 owner, in step 206. Alternatively, the new owner can be requested to obtain 115 an
29 object copy from the originating object server, e.g., via the Internet 110. Those skilled
30 in the art will appreciate that the load checking step 204 can be performed proactively,
i.e., periodically or in response to an identified overload or overload trend 1021 due at
least in part to a high forwarding frequency 1011--for a given object id/partition id
1010 and cache server (ownership 1012). If so, then in step 205, the load monitor finds
an underloaded (or less loaded) cache server, assigns it as the new (or shared) owner
of the requested object, and possibly sends a copy of the object to the new (or shared)
owner as above. Conversely, if a shared ownership model is used, in step 208, when
the load condition 10211 and forwarding frequency 10111 for a shared ownership
object (p 10101) drops below a predetermined threshold, in step 209, the shared
ownership (B, A 10121) can be merged to a single ownership and one of the copies

purged from one of the cache servers A 10121, e.g., to make room for another hot object. (Emphasis added, Jordan, column 6, line 50-column 7, line 22).

As evident from the underlined portion above, Jordan explains how the load monitor receives a forward request 125 from a cache server 150 because of a cache miss, and as provided in a step 201, the load monitor checks to see if the requested object is in a caching table. If the requested object is not in the caching table, in addition to updating its forwarding frequency, the load monitor then finds an underloaded or less loaded cache server and assigns it as the new or shared owner of the requested object (Jordan, column 6, lines 63-65). Thus, the *ownership information of the object in the caching table is updated* (Jordan, column 6, lines 65-67). In Jordan, it is these steps that update the caching table, which identifies a less loaded cooperating cache server. But, it is important to note that the identification takes place in the caching table and not in an equivalent of applicants' recited *intake message*.

Jordan also teaches that a request can be forwarded to the new owner, as shown in a step 206. But again, a forwarded request does not identify a less-loaded server. Instead, the caching table identifies the owner or less loaded server. Accordingly, the request 125 (i.e., forwarded requests 125) is also not equivalent to applicants' *intake message* because of this lack of identification in the request.

Based upon the preceding comments, it should be apparent that the rejection of independent Claim 1 over O’Neil in view of Jordan should be withdrawn, because O’Neil and Jordan do not teach or suggest all of the elements of independent Claim 1. Specifically, the cited art does not teach or suggest any equivalent to an intake message or describe a server that serves as a new intake, as recited by applicants’ amended claim.

Claims 2-18 ultimately depend from independent Claim 1. Because dependent claims inherently include all of the steps or elements of the independent claims from which the dependent claims ultimately depend, dependent Claims 2-18 are patentable for at least the same reasons discussed above with regard to independent Claim 1. Therefore, the rejection of dependent Claims 2-18 under 35 U.S.C. § 103(a) over O’Neil and in view of Jordan should be withdrawn.

Rejection of Independent Claim 19

Independent Claim 19 is directed to a system for distributing a processing load in a cluster. Clearly, for the same reasons already noted above in regard to independent Claim 1, Claim 19 also distinguishes

1 over O'Neil and Jordan, because O'Neil and Jordan do not teach or suggest providing an *intake* message
2 from a first resource to a plurality of resources in the cluster, for *identifying* the second resource as the new
3 intake, wherein the new intake is a single entry-point service resource in the cluster.

4 Accordingly, the rejection of independent Claim 19 over O'Neil and further in view of Jordan
5 should be withdrawn for the reasons discussed above. Because dependent claims include all of the elements
6 of the independent claim from which the dependent claims ultimately depend and because O'Neil and
7 Jordan do not teach or suggest all of the elements of independent Claim 19, the rejection of dependent
8 Claims 20-23, under 35 U.S.C. § 103(a) over O'Neil and Jordan should also be withdrawn, for at least the
9 same reasons as the rejection of Claim 19.

10 Rejection of Independent Claim 24

11 Independent Claim 24 is directed to a method of distributing a processing load among a cluster of
12 nodes, each node providing at least one of a plurality of different types of services. Clearly, for the same
13 reasons already noted above in regard to independent Claim 1, this claim also distinguishes over O'Neil and
14 Jordan, because O'Neil and Jordan do not teach or suggest providing an *intake* message from the first node
15 designated as the intake for the first instance, to the nodes in the cluster *identifying* the second instance as the
16 new intake for the first type of service, wherein the new intake is a *single entry-point service resource in the*
17 *cluster*.

18 Accordingly, the rejection of independent Claim 24 under 35 U.S.C. § 103(a) over O'Neil and
19 Jordan should be withdrawn based on the reasons discussed above. Because dependent claims include all of
20 the elements of the independent claim from which the dependent claims ultimately depend, and because
21 O'Neil and Jordan do not teach or suggest all of the elements of independent Claim 24, the rejection of
22 dependent Claims 25-35, under 35 U.S.C. § 103(a) over O'Neil and Jordan should also be withdrawn, for at
23 least the same reasons as the rejection of Claim 24.

24 Discussion of the Rejection of Independent Claim 36

25 Independent Claim 36 is directed to a system for distributing a processing load in a cluster of
26 resources. However, O'Neil and Jordan fail to teach or suggest providing an *intake* message from a first
27 resource to a plurality of resources in a cluster *identifying* a second resource as an intake, wherein the intake
28 is a single entry-point service resource in the cluster. For the same reasons already noted above in regard to
29 independent Claim 1, this claim also distinguishes over O'Neil and Jordan. Accordingly, the rejection of
30 independent Claim 36 under 35 U.S.C. § 103(a) over O'Neil and Jordan should be withdrawn.

In view of the Remarks set forth above, it will be apparent that the claims in this application define a novel and non-obvious invention. The application is in condition for allowance and should be passed to issue without further delay. Should any further questions remain, the Examiner is invited to telephone applicants' attorney at the number listed below.

Respectfully submitted,

/sabrina macintyre/
Sabrina K. MacIntyre
Registration No. 56,912

SKM/RMA:elm